



(19)

Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 1 049 299 A1

(12)

## EUROPEAN PATENT APPLICATION

(43) Date of publication:  
02.11.2000 Bulletin 2000/44

(51) Int. Cl.<sup>7</sup>: H04L 12/56

(21) Application number: 00303491.5

(22) Date of filing: 26.04.2000

(84) Designated Contracting States:  
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU  
MC NL PT SE  
Designated Extension States:  
AL LT LV MK RO SI

(30) Priority: 30.04.1999 US 302375

(71) Applicant:  
Nortel Networks Limited  
Montreal, Quebec H2Y 3Y4 (CA)

(72) Inventors:  
• Chapman, Alan Stanley John  
Kanata, Ontario K2K 1V5 (CA)  
• Kung, Hsiang-Tsung  
Lexington, Massachusetts 02173 (US)

(74) Representative:  
Ertl, Nicholas Justin  
Elkington and Fife,  
Prospect House,  
8 Pembroke Road  
Sevenoaks, Kent TN13 1XR (GB)

### (54) Method and apparatus for bandwidth management of aggregate data flows

(57) The present invention relates to a transmission device implementing a flow control mechanism for aggregate trunks. The transmission device can be implemented as a router that includes an input for receiving aggregate traffic streams, an output for releasing the aggregate traffic streams to a destination point and a control unit capable to regulate the rate of release of packets from the output. Specifically, the flow control operation effected by the control unit is dependent on receipt of acknowledgement messages issued at the

destination point, each acknowledgement message confirming the receipt of one or more particular packets at the destination point. The control unit will continuously increase the packet sending rate until a packet is lost in the network between the transmission device and the destination point. On detection of packet loss, based on the lack of one or more corresponding acknowledgement messages from the destination point, the control unit will reduce the packet sending rate.




Figure 2

EP 1 049 299 A1

**Description**

**[0001]** The present invention relates to the field of digital data transmission. More specifically, it pertains to a method and apparatus for performing traffic management functions, such as bandwidth management of aggregate traffic streams between two points of a data transport network.

**[0002]** A typical data transport network operates in a connectionless mode whereby there is no negotiation between the transmitter/receiver and the network with regard to the type or quantity of traffic that is to be sent. The transmitter simply sends the traffic on the network, and relies on the network components to deliver that traffic to the receiver accurately. These network components consist typically of routing nodes (also known as routers or switches) joined by physical links. The main function of the routing nodes is to direct incoming packets to the appropriate outgoing links.

**[0003]** Many communication network applications have adopted the Internet Protocol (IP), a library of routines called on by various network communications applications for transporting packets of data from node to node. The Transmission Control Protocol (TCP) is used for most IP network transactions and provides for both reliable transmission and rate adaptation in order to adjust to available network bandwidth. TCP allows adaptive use of available bandwidth and accommodation of different rates at different points in the internetwork or network.

**[0004]** Within the IP networks, there is a growing interest in managing data traffic in the form of aggregate flows. A flow, also referred to as a connection, is a stream of traffic under control of one sender within the network. Specifically, aggregating individual traffic flows into one larger flow simplifies the management of the network; routes can be pinned down for the aggregates more simply than if all flows were visible. Treatments such as bandwidth allocation or forwarding priority can be assigned to the aggregate rather than having to manage it on a per-flow basis. In the future, it may be desirable to provide flow control for these aggregates, similar to TCP, so that they can compete for available bandwidth in a controlled manner.

**[0005]** In Internet terminology, aggregating data traffic by encapsulating them into a single IP packet stream is often called tunneling, where the encapsulated traffic flow through the IP network is referred to as an aggregate trunk or tunnel. When the number of aggregate trunks becomes very large, it may be desirable to add another level of aggregation by putting several such trunks into a new, larger trunk between chosen points in the network. One way to do this is to encapsulate the packets again with another TCP/IP header. Unfortunately, adding more headers to data packets is inefficient in bandwidth use and may be problematical if the longer packets violate Maximum Transport Unit (MTU) limits. Multi-protocol Label Switching

(MPLS), another form of aggregation currently under development, minimizes the overhead by using a very small label instead of a full TCP/IP header. Unfortunately, the small header envisaged for MPLS does not provide room for the information usually carried in a TCP/IP flow controlled connection.

**[0006]** The background information herein clearly shows that there exists a need in the industry to provide an improved mechanism for ensuring flow control of an aggregate stream between two points in a network.

**[0007]** According to the invention, there is provided a transmission device for forwarding aggregate traffic streams towards a destination point, an aggregate traffic stream being comprised of a plurality of packets, said transmission device comprising:

an input for receiving the aggregate traffic streams; an output for forwarding the aggregate traffic streams to the destination point; a control unit having an input for receiving a message issued from the destination point allowing said transmission device to determine if data relating to a certain packet released from said output has been received at the destination point, said control unit being operative to regulate a rate at which packets are released from said output at least in part in dependence upon said message.

**[0008]** The transmission device can detect and react to congestion within the network and allows for rate regulation of the aggregate flow without adding marking data to every existing packet.

**[0009]** According to one embodiment of the present invention, the forwarding point is analogous to a virtual sender, implemented by a router. Similarly, the destination point is analogous to a virtual receiver, also implemented by a router. The two routers are connected across a network of components by an aggregate trunk. The forwarding router includes a control unit responsible for implementing the flow control of the aggregate trunk. The control unit observes incoming packets to find for each packet a packet identifier. More specifically, the sender will look at the packet structure to find a bit sequence that is sufficiently unique to distinguish the packet from other packets in the stream. This packet identifier is recorded in a data structure. The control unit processes this data structure in conjunction with successive acknowledgement messages received from the virtual receiver, also containing the packet identifier (read or observed by the virtual receiver and included in the acknowledgement message), to determine if packets are being dropped. In the affirmative, the packet forwarding rate is decreased, otherwise it is increased.

**[0010]** According to yet another embodiment of the present invention, control packets are inserted into the aggregate trunk traffic stream at the forwarding point. The destination point is only responsible for acknowledging the control packets, and not the data packets of

the aggregate traffic stream. Thus, flow control of the aggregate trunk can be implemented by looking only at the control packets. The control packets can be distinguished from the other packets based on some arbitrary identifier that is present in the control packets but absent in the other (non-control) packets. A possible variant is to insert packet sequence information in the control packets, such that the destination point can send a control message to the forwarding point only when a control packet is missing.

[0011] In yet another embodiment, some of the aggregate traffic stream packets are marked at the forwarding point, and the destination point only acknowledges these marked packets. A possible variant is to use the marking data to convey packet sequence information, such that the destination point can send a control message to the forwarding point only when a marked packet is missing.

[0012] The Invention also provides a method for controlling the flow of an aggregate traffic stream between a transmission device and a destination point, the aggregate traffic stream being comprised of a plurality of packets, said transmission device comprising:

an input for receiving the aggregate traffic streams; an output for forwarding the aggregate traffic streams to the destination point; said method comprising regulating a rate at which packets are released from said output at least in part in dependence upon a message issued at the destination point allowing the transmission device to determine if data relating to at least one packet released from said output has been received at the destination point.

[0013] Examples of the invention will now be described in detail with reference to the accompanying drawings, in which:

Figure 1 is a block diagram of a flow controlled aggregate stream in an IP network, in accordance with an embodiment of the present invention; Figure 2 is a block diagram of a router (virtual sender) shown in Figure 1; Figure 3 illustrates a TCP congestion window function used in an embodiment of the present invention; Figure 4 is a flowchart illustrating the operation of a program element in the router depicted in Figures 1 and 2, which implements the flow control of the aggregate stream.

[0014] A conventional IP network implements bandwidth sharing among host machines using the Transport Control Protocol (TCP). Although data flow in the network can be bi-directional, it is usual to refer to the originator of a particular piece of data as the sender and the other end as the receiver. In TCP, the sender (sender

host machine) constantly tests the network to see if more bandwidth is available and uses the loss of a packet determined by sequence numbers of TCP packets as an indication to decrease its rate. Any lost packets are sent again so that there is a reliable flow of traffic. The loss of too many packets can cause the TCP connection to enter the timed out state.

[0015] Each packet contains a sequence number, which increases according to the number of bytes transmitted. The receiver acknowledges packets using this numbering scheme and always acknowledges the latest packet received in correct sequence. It may acknowledge each packet individually or wait in order to reduce overhead. It should definitely send an acknowledgement at least every second packet. Such acknowledgements can be piggybacked on data packets going in the reverse direction (from receiver to sender), or can be sent as stand-alone packets with no data. If a packet is received which is not in correct sequence, the receiver will immediately send an acknowledgement but the sequence number it acknowledges will be that of the last packet which was received in correct sequence. It should be noted that the sequence number in a packet corresponds to the last byte in the packet and the acknowledgement contains the next expected in-sequence byte number and thus acknowledges all bytes up to that number. In general terminology, a packet is acknowledged when the receiver reports that the next expected byte number is later than any bytes contained in that packet.

[0016] The general characteristic of TCP is that it is self-clocking. That is to say, the sender will wait for an acknowledgement from the receiver for the packets already sent before sending more packets. If the sender waited for each individual packet to be acknowledged than the maximum rate that the connection could achieve would be one packet per round trip time of the connection. To increase the sending rate while keeping the self-clocking nature of the protocol, the sender is allowed to send some number of packets while waiting for an earlier packet to be acknowledged. This number of packets is called the window. The receiver itself may constrain the size of the window in order to limit its buffer requirement.

[0017] The current size of the window is called the congestion window and can vary between one packet and the maximum that the receiver is prepared to accept. As the sender receives acknowledgements, the window slides forward and also increases in size. An increase in size allows the connection to run faster. If a packet is not acknowledged it will eventually be considered lost and this loss is assumed to be a result of congestion at some merge point. The sender, in addition to re-transmitting the packet, will reduce the window size. The slow and gradual increase in window size then begins again.

[0018] One embodiment of the present invention uses a mechanism to manage the data flow rate over a

virtual connection between two points within a network that has some common points to the TCP process. This mechanism is referred to as *TCP-like* since the virtual sender receives an indication that forwarded packets have successfully passed through the virtual receiver, and the virtual sender uses the presence or absence of these acknowledgements to reduce and increase its rate. The term *virtual* refers to the fact that the connection corresponds to a control loop running between two points in the network (not necessarily host machines or encapsulation/decapsulation points). Further, the traffic flowing over the virtual connection between virtual sender and virtual receiver is not being generated at the virtual sender.

[0019] Figure 1 illustrates an embodiment of a flow controlled aggregate connection (trunk), hereafter referred to as the virtual connection 106. The virtual connection 106 can be initiated between any two points in the network. Specific to this example, one of these points is referred to as the virtual sender, implemented as a router 100, while the other is referred to as the virtual receiver, implemented as router 102. Router 100 is the forwarding point and determines the packet sending rate (i.e. rate at which packets are forwarded onward). Router 102 is the destination point and monitors and acknowledges all packets received from the virtual sender 100. The functionality of these routers will be described in further detail below. The connection between virtual sender 100 and virtual receiver 102 may comprise a plurality of network 104 components, such as other routers.

[0020] As shown in Figure 2, router 100 (virtual sender) generally includes a control unit 222 and an internal bus 204. The control unit 200 itself includes a memory 200 and a processor 202, and is responsible for implementing the flow control of the aggregate trunk 106, as will be described in further detail below. The internal bus 204 interconnects the router 100 components, enabling data and control signals to be exchanged between them. Links 224 and 226 provide for the input and output of data to and from the components of the control unit 222, in cooperation with the internal bus 204. Input ports 206, 208 and 210 connect the router 100 to physical links 216, 218 and 220, respectively, for receiving network traffic over the different link flows, where these flows have all been assigned (as per customer choice) to the virtual trunk initiated between virtual sender 100 and virtual receiver 102. Virtual port 212 is an output port that connects the router 100 to the virtual connection 106 for forwarding aggregate traffic streams to the virtual receiver 102.

[0021] The memory 200 of the control unit 222 includes a First-In-First-Out (FIFO) buffer 214 that can hold data packets received at the various input ports. The purpose of the buffer is to provide a temporary storage mechanism for holding all traffic allocated to the virtual connection 106 (i.e. traffic from all three input ports 206, 208 and 210) until a decision is made by the con-

5 trol unit 222 to forward the packets onward over the virtual connection 106. The physical configuration of the buffer 214 does not need to be described in detail because such components are readily available in the marketplace and the selection of the appropriate buffer mechanism suitable for use in the present invention is well within the reach of a person skilled in the art. The memory 200 also supports a standard TCP congestion window and a storage data structure, the latter including a packet identity table and a packet scoreboard table, for use by the virtual sender 100 control unit 222, all three of which will be described in further detail below.

10 [0022] The memory 200 further contains a program element that regulates the flow control operation of the router 100. That program element is comprised of individual instructions that are executed by the processor 202, for implementing the flow control of the aggregate connection 106 between virtual sender 100 and virtual receiver 102. This program element will be described in further detail below.

15 [0023] In this embodiment the virtual sender 100 adds no information to the data packets that it receives and forwards on to the virtual receiver 102. The sender 100 and the receiver 102 achieve packet identification by inspecting a part or all of the packet's existing contents. This part is called "packet identifier". This means that to perform packet identification, say at the sender 100, the sender 100 will look at the packet structure to find a bit sequence that is sufficiently unique such as to

20 allow to distinguish this packet from other packets in the flow. Once a packet is received by the virtual receiver 102, the receiver 102 reads some or all of the contents of the packet in order to determine its unique identity (packet identifier) in the same manner as the sender 100. The sender 100 and the receiver 102 are set to look at the same part of the packet when they perform packet identification, so that they extract the same bit sequence. The particular packet part or section looked at by the virtual sender 100 and the receiver 102 can be arbitrarily set. It should also be noted that the entire packet may serve as the packet identifier, as the packet identifier is not necessarily limited to a portion of the packet.

25 [0024] The receiver 102 then sends an *Acknowledgement Message* containing this packet identifier to the sender 100 in order to confirm receipt of the particular packet. Assume for this example that the packets being forwarded from the virtual sender 100 to the virtual receiver 102 originated at a TCP/IP encapsulation point. Each packet will contain a unique sequence number for the encapsulated TCP connection and an identifier for the trunk (for instance the source and destination IP addresses), forming for each packet the requisite packet identifier. Thus, in this case, the sender 100 and the receiver 102 are programmed to look at the sequence number field to generate a packet identifier.

30 [0025] In another example, the packets may be coming from a TCP sender in a host, where typically

there is already a well understood way of using IP addresses, TCP port numbers and a sequence number to get a unique packet identifier. In yet another example, the packets may be unencapsulated Unacknowledged Datagram Protocol (UDP) packets, MPLS packets or encrypted packets, for which a hash coded form of the packet contents can be used to form the unique packet identifier. In such a case, it is necessary to take a large enough number of bytes from the packet in order to ensure that the hash code generated has a high probability of being unique. Further, the bytes selected from the packet must not include fields that might change between virtual sender 100 and virtual receiver 102, such as a Time-To-Live (TTL) indication. Note that over an MPLS path the packet is not changed, so there is more freedom in choosing the bytes to use.

[0026] Most preferably, the virtual receiver 102 acknowledges every packet received from the virtual sender 100. Note that the virtual receiver 102 can compose *Acknowledgement Messages* containing one or more packet receipt acknowledgements and insert them into the flow of traffic traveling towards the virtual sender 100.

[0027] The router 102 (virtual receiver) is structurally very similar to the virtual sender 100, and as such will not be described from a structural point of view. Although for this example the virtual receiver 102 as shown in Figure 1 is the destination point for only one trunk, notably virtual connection 106, in a typical telecommunications network such a router may act as the terminating point for a plurality of such trunks, where each trunk forms a particular virtual connection between the virtual receiver 102 and a different virtual sender. Since the virtual receiver 102 has no knowledge of which packets come from which sender, it acknowledges receipt of a packet to the virtual sender for each of the trunks it is terminating. An *Acknowledgement Message* is broadcast by the virtual receiver 102 to all of the virtual senders to which it is connected. Note that in the case of an MPLS packet stream, the labels give a direct indication of which sender to respond to, avoiding the requirement for an acknowledgement broadcast. For the sake of the present example, assume hereinafter that the virtual receiver 102 is connected to only one sender, notably virtual sender 100.

[0028] Since there is no sequence number information added to the packets passing through virtual sender 100 with respect to the virtual connection 106, later *Acknowledgement Messages* sent by the virtual receiver 102 can not imply receipt of earlier packets, as is currently implemented by the standard TCP acknowledgement protocol. The loss of such an *Acknowledgement Message* may have an undesirable effect by suggesting the loss of many packets. In the context of the present invention, it is preferred that the virtual receiver 102 repeats an *Acknowledgement Message* until virtual sender 100 acknowledges receipt of it. Under this implementation, the virtual sender 100

acknowledges every *Acknowledgement Message* received from the virtual receiver 102 with a *Return Acknowledgement Message*.

[0029] When the virtual sender 100 transmits a packet over the virtual connection 106, it first determines the packet's unique identifier and its length, with which it calculates an equivalent byte sequence number. It then stores this sequence number by mapping it to the particular packet identifier in the aforementioned packet identity table of memory 200. When the virtual sender 100 receives *Acknowledgement Messages* from the virtual receiver 102, it reconstructs sequence number information by extracting the packet identifier from the *Acknowledgement Message* and reading the corresponding sequence number from the packet identity table.

[0030] Byte sequence numbers manage the congestion window stored in memory 200 of the virtual sender 100, as per standard TCP protocol. A typical example of such a congestion window is that shown in Figure 3, where the left-hand side of the congestion window 300 is aligned with the earliest packet that has been sent but not acknowledged. Since the congestion window 300 is a feature of TCP protocol that is well known in the art and has been well documented, it will not be described in further detail. The control unit 222 of the virtual sender 100 keeps a history of identities of packets passed forward but not yet acknowledged and processes this history in conjunction with arriving *Acknowledgement Messages* to adjust the pointers to the congestion window 300 and thus regulate the rate at which packets are forwarded over the virtual connection 106. In a particular example, this history is stored in a packet scoreboard in memory 200 of the virtual sender 100. The packet scoreboard is an array of binary variables, indexed by the packet byte sequence numbers. When a packet is sent over the virtual connection 106, the variable indexed by the packet identifier's corresponding sequence number is set to 1. When a packet is acknowledged, the variable indexed by the packet identifier's corresponding sequence number is cleared to 0 and, if it is the earliest packet, the left side of the congestion window 300 is moved to align with the earliest packet sent but not yet acknowledged. Thus, the packet scoreboard is a rolling list of packets that corresponds to the current congestion window 300, allowing the virtual sender 100 to detect the equivalent of multiple duplicate acknowledgements and decide whether a packet has been lost. Normal TCP mechanisms are used by the virtual sender 100 to implement window increase and decrease for controlling the sending rate of packets being transmitted over the virtual connection 106. Specifically, the virtual sender 100 will progressively increase the congestion window 300 size, and thus the sending rate, until one or more packets are lost in the network between itself and the virtual receiver 102. The virtual sender 100 detects packet loss in dependence upon the *Acknowledgment Messages*, or

lack thereof, received from the virtual receiver 102. On detecting packet loss (i.e. a packet forwarded over the virtual connection 106 to the virtual receiver 102 is never acknowledged), the virtual sender 100 reduces its congestion window 300 in order to decrease the sending rate. A lower sending rate on the aggregate trunk will be reflected as packet discard at the FIFO buffer 214, which in turn will eventually be seen by the real sender.

[0031] Note that once a packet loss has been acted on by the virtual sender 100, the packet is marked as delivered. The virtual receiver 102 does not wait for resends of packets that have been lost, but rather simply passes on all packets as fast as possible to avoid adding to network latency. If required, the real sender, that is the originator of the packet (i.e. TCP sender in a host, TCP/IP encapsulation point, etc), will achieve packet reliability. Consequently, the virtual sender 100 does not retransmit lost packets or store them once they are sent over the virtual connection 106.

[0032] Figure 4 provides a complete flowchart illustrating an example of the operation of the program element stored in the memory 200, and executed by the processor 202, that regulates the operation of the virtual sender 100, in particular the flow control of the aggregate trunk 106. At step 400, a packet from a traffic flow assigned to the aggregate trunk 106 arrives at input port 206 of the virtual sender 100, and is stored in FIFO buffer 214. At step 402, a unique packet identifier is determined and a corresponding byte sequence number calculated for the packet. The two values are stored and mapped to each other in the packet identity table in memory 200 of the virtual sender 100. The packet is next sent over the aggregate trunk 106 to the virtual receiver 102, at step 404. Immediately following transmission, at step 406, the packet scoreboard in memory 200 of the virtual sender 100 is accessed and the variable indexed by the packet's sequence number is set to 1, indicating that the packet has been sent but has not yet been acknowledged. The program next waits for a packet *Acknowledgement Message* from the virtual receiver 102 at step 408. When such a message is received, the packet identifier is read from the message and the corresponding byte sequence number obtained from the packet identity table, at step 410. Assume for this example, that the message received acknowledges receipt of the particular packet sent out earlier at step 404. At step 412, the virtual sender 100 sends a *Return Acknowledgement Message* back to the virtual receiver 102 confirming receipt of the particular packet *Acknowledgement Message*, where this *Return Acknowledgement Message* includes the particular packet identifier. Next, the appropriate variable in the packet scoreboard is cleared, indicating that the packet has been sent and acknowledged, at step 414. At step 416, the program checks to determine whether the acknowledged packet corresponds to the earliest packet in the congestion window 300. If so, the window is shifted at step 418 to align with the earliest packet

that has been sent but not acknowledged.

[0033] In an alternative embodiment of the present invention, simple marking of the packets is implemented by a marking unit at the virtual sender, allowing for certain packets to be selectively marked for acknowledgement by the virtual receiver. Thus flow control is implemented by looking only at the marked packets, which is somewhat analogous to *sampling* the network between virtual sender and receiver in order to provide the virtual sender with a view of the existing congestion conditions. Such an implementation reduces the processing power required as well as the amount of acknowledgement traffic because fewer packets need to be processed. Objectively, more bandwidth is consumed since marking data must be inserted in some of the packets. The marking data is a unique combination of bits placed by the marking unit at a known location in the packet. Such a location can be in an existing field of the packet or a new field can be created and appended to the packet. The receiver 102 is set to look for this unique combination of bits at the predetermined location and examines each packet for this feature. Only packets containing this feature are being acknowledged.

[0034] In a possible variant, the marking data in addition to being an identifier may also include a sequence number. Thus the marking data may be structured as a two-element data block, the first element indicating that this is a marked packet and the second element indicating the sequence number of the packet that is different from one packet to the other. The sequence number is useful because it allows the virtual receiver to determine which packets have been dropped only by observing the sequence continuity in the received marked packets. Thus the virtual receiver may issue a control message to indicate the loss of one or more marked packets, rather than issuing acknowledgement messages to confirm the receipt of each marked packet. This is preferred since confirming loss of marked packets will require fewer messages from the virtual receiver to the virtual sender than confirming receipt of all marked packets.

[0035] In yet another alternative embodiment, the virtual sender includes a packet generator unit for generating control packets. These control packets are inserted into the aggregate trunk traffic stream and flow control implemented by looking only at the control packets, thus eliminating the requirement for robust identification of packets. Although a control packet could be sent for each user packet, it would be more efficient for control packets to be sent only as often as required to maintain a responsive control loop. The control packets are *sampling* the network between virtual sender and virtual receiver in order to give the sender a view of the existing congestion conditions. In such an implementation of a virtual TCP connection, the virtual sender maintains accounting on the whole traffic flow in order to perform window management. Although additional

bandwidth is required to transport the control packets, this embodiment uses less bandwidth than prior art approaches where encapsulation requires that every data packet contain control information.

[0036] In a possible variant, the control packets may include sequence number information, removing the need for identification of user packets. Further, the virtual receiver may send a control message to notify the virtual sender of one or more missing packets, as opposed to sending acknowledgement messages upon receipt of every control packet. Such an alternative reduces the processing power required at a router (no packet identification required) and reduces acknowledgement traffic at the cost of some bandwidth.

[0037] The above description of a preferred embodiment under the present invention should not be read in a limitative manner as refinements and variations are possible without departing from the spirit of the invention. The scope of the invention is defined in the appended claims and their equivalents.

#### Claims

1. A transmission device for forwarding aggregate traffic streams towards a destination point, an aggregate traffic stream being comprised of a plurality of packets, said transmission device comprising:

an input for receiving the aggregate traffic streams;  
an output for forwarding the aggregate traffic streams to the destination point;  
a control unit having an input for receiving a message issued from the destination point allowing said transmission device to determine if data relating to a certain packet released from said output has been received at the destination point, said control unit being operative to regulate a rate at which packets are released from said output at least in part in dependence upon said message.

2. A transmission device as defined in claim 1, wherein the packets include respective identifiers allowing to distinguish one packet from another.

3. A transmission device as defined in claim 1 or 2, wherein the message is an acknowledgement message for notifying the transmission device that a packet has been received.

4. A transmission device as defined in any preceding claim, further comprising a marking unit for selectively marking certain packets of the aggregate traffic streams received at said input with marking data prior to their release from said output, said marking data allowing to distinguish the marked packet from another packet in the traffic stream, and wherein

the message allows the transmission device to determine if a certain marked packet released from said output has been received at the destination point.

- 5. A transmission device as defined in claim 4, wherein said message contains the marked data of a packet.
- 10. 6. A transmission device as defined in claim 1, 2 or 3, wherein said control unit forwards the aggregate traffic stream to the destination point without adding data elements to the packets of the aggregate traffic streams.
- 15. 7. A transmission device as defined in claim 6, wherein said message conveys information relative to a packet identifier.
- 20. 8. A transmission device as defined in claim 7, wherein said control unit includes a data structure, said control unit being operative for recording the identifiers of packets released at said output for forwarding to the destination point in said data structure, said control unit being operative to process said data structure in conjunction with successive messages received from the destination point to regulate the rate at which packets are released from said output.
- 25. 9. A transmission device as defined in claim 8, wherein said control unit is operative to process said data structure in conjunction with successive messages received from the destination point in order to determine whether packets forwarded to the destination point have not been received at the destination point.
- 30. 10. A transmission device as defined in claim 9, wherein if at least one packet has not been received at the destination point, said control unit is operative to reduce a rate of release of the packets from said output.
- 35. 11. A transmission device as defined in claim 10, wherein said control unit is operative to progressively increase a rate of release of the packets from said output until a packet is not received at the destination point.
- 40. 12. A transmission device as defined in claim 1, further comprising a packet generator unit for generating and inserting in the aggregate traffic stream control packets, and wherein the output is for releasing the aggregate traffic streams and the control packets inserted in the aggregate traffic stream to the destination point, and wherein the message issued from the destination point allows the transmission device
- 45. 50. 55.

to determine if a control packet released from said output has been received at the destination point, the control packet comprising the data relating to a packet.

13. A transmission device as defined in claim 12, wherein each control packet includes sequence information indicative of a sequence of insertion of the control packet relative to other control packets inserted by the packet generator in the aggregate traffic stream. 10

14. A transmission device as defined in claim 13, wherein said packet generator inserts control packets into the aggregate traffic streams at predetermined intervals. 15

15. A transmission device as defined in claim 14, wherein said control message is an acknowledgement message that conveys to the transmission device the sequence information of at least one control packet received at the destination point.

16. A transmission device as defined in claim 15, wherein said control unit is operative to process successive acknowledgement messages received from the destination point in order to determine whether control packets forwarded to the destination point have not been received at the destination point. 25 30

17. A transmission device as defined in claim 16, wherein if at least one control packet has not been received at the destination point, said control unit is operative to reduce a rate of release of the data packets from said output. 35

18. A transmission device as defined in claim 17, wherein said control unit is operative to progressively increase a rate of release of the data packets from said output until a control packet is not received at the destination point. 40

19. A method for controlling the flow of an aggregate traffic stream between a transmission device and a destination point, the aggregate traffic stream being comprised of a plurality of packets, said transmission device comprising: 45

an input for receiving the aggregate traffic streams;  
an output for forwarding the aggregate traffic streams to the destination point;  
said method comprising regulating a rate at which packets are released from said output at least in part in dependence upon a message issued at the destination point allowing the transmission device to determine if data relat- 50 55

ing to at least one packet released from said output has been received at the destination point.

5 20. A method as defined in claim 19, in which the packets have respective identifiers allowing to distinguish one packet from another

21. A method as defined in claim 19 or 20, wherein the message is an acknowledgement message for notifying the transmission device that a packet has been received.

22. A method as defined in any one of claims 19 to 21, wherein the regulation of the rate at which packets are released from said output is performed without adding any data to packets released from said output. 20

23. A method as defined in claim 22, wherein said message conveys information relative to a packet identifier.

24. A method as defined in claim 23, comprising recording the identifiers of packets released at said output for forwarding to the destination point in a data structure, and processing said data structure in conjunction with successive messages received from the destination point to regulate the rate at which packets are released from said output. 25 30

25. A method as defined in claim 24, comprising processing said data structure in conjunction with successive messages received from the destination point in order to determine whether packets forwarded to the destination point have not been received at the destination point. 35

26. A method as defined in claim 25, wherein if at least one packet has not been received at the destination point, said method comprises reducing a rate of release of the packets from said output. 40

27. A method as defined in claim 26, wherein said method comprises the step of progressively increasing a rate of release of the packets from said output until a packet is not received at the destination point. 45

28. A method as defined in claim 19, the method comprising inserting control packets into the aggregate traffic stream, and wherein the message issued from the destination point allows the transmission device to determine if a control packet released from said output has been received at the destination point, the control packet comprising the data relating to the packet. 50 55

29. A method as defined in claim 28, wherein each control packet includes sequence information indicative of a sequence of insertion of the control packet relative to other control packets inserted in the aggregate traffic stream. 5

30. A method as defined in claim 29, wherein said control message is an acknowledgement message representative of the sequence information of at least one control packet received at the destination point. 10

31. A method as defined in claim 30, wherein said method comprises the step of processing successive acknowledgement messages received from the destination point in order to determine whether control packets have not been received at the destination point. 15

32. A method as defined in claim 31, wherein if at least one control packet has not been received at the destination point, said method comprises the step of reducing a rate of release of data packets from said output. 20

33. A method as defined in claim 32, wherein said method comprises the step of progressively increasing a rate of release of the data packets from said output until a control packet is not received at the destination point. 25

34. A method as defined in claim 20, wherein the method comprises selectively marking certain packets of the aggregate traffic streams received at said input with marking data prior to their release from said output, the marking data allowing to distinguish one marked packet from another marked packet, and wherein the message allows the transmission device to determine if a certain marked packet released from said output has been received at the destination point. 30 35 40

45

50

55




Figure 1




Figure 2




Figure 3




Figure 4



European Patent  
Office

## EUROPEAN SEARCH REPORT

Application Number  
EP 00 30 3491

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
X	<p>PAZOS C M ET AL: "USING BACK-PRESSURE TO IMPROVE TCP PERFORMANCE WITH MANY FLOWS" PROCEEDINGS IEEE INFOCOM. THE CONFERENCE ON COMPUTER COMMUNICATIONS, US, NEW YORK, NY: IEEE, 21 March 1999 (1999-03-21), pages 431-438, XP000870717 ISBN: 0-7803-5418-4 * the whole document *</p> <p>---</p>	1,2,6, 19,20,22	H04L12/56
X	<p>CHANG K: "IP Layer Per-flow Queueing and Credit Flow Control" PH.D. DISSERTATION, [Online] January 1998 (1998-01), pages 18-28, XP002146481 Harvard University Retrieved from the Internet: &lt;URL: http://www.eecs.harvard.edu/~chang/papers/thesis.ps&gt; [retrieved on 2000-08-17] * page 18 - page 28 *</p> <p>---</p>	1,2,4,5, 19,20,22	
X	<p>EP 0 712 220 A (IBM) 15 May 1996 (1996-05-15)</p> <p>* page 4, line 3 - line 31 * * page 4, line 53 - page 6, line 5 * * page 6, line 47 - page 7, line 14 * * figures 4,5,7,8,10 *</p> <p>---</p>	1,2,6-8, 19,20, 22-24	
A		3-5, 9-18,21, 25-34	H04L
A	<p>EP 0 596 159 A (BELL TELEPHONE MFG ;ALCATEL NV (NL)) 11 May 1994 (1994-05-11)</p> <p>* abstract *</p> <p>---</p> <p>-/-</p>	1-34	
<p>The present search report has been drawn up for all claims</p>			
Place of search	Date of completion of the search	Examiner	
BERLIN	4 September 2000	Eraso Helguera, J	
CATEGORY OF CITED DOCUMENTS		<p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons</p> <p>&amp; : member of the same patent family, corresponding document</p>	
<p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p>			



DOCUMENTS CONSIDERED TO BE RELEVANT			CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	
A	CHEN M -S ET AL: "HOP-BY-HOP FLOW CONTROL ON UNRELIABLE CONNECTIONS" INTERNATIONAL CONFERENCE ON COMMUNICATIONS, US, NEW YORK, IEEE, vol. -, 7 July 1987 (1987-07-07), pages 1683-1687, XP002036563 * the whole document *	1-34	
A	OZVEREN C M ET AL: "RELIABLE AND EFFICIENT HOP-BY-HOP FLOW CONTROL" IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS, US, IEEE INC. NEW YORK, vol. 13, no. 4, 1 May 1995 (1995-05-01), pages 642-650, XP000501260 ISSN: 0733-8716 * abstract *	1-34	
			TECHNICAL FIELDS SEARCHED (Int.Cl.7)
The present search report has been drawn up for all claims			
Place of search	Date of completion of the search		Examiner
BERLIN	4 September 2000		Eraso Helguera, J
CATEGORY OF CITED DOCUMENTS			
X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document			
T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document			

**ANNEX TO THE EUROPEAN SEARCH REPORT  
ON EUROPEAN PATENT APPLICATION NO.**

EP 00 30 3491

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The members are as contained in the European Patent Office EDP file on The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

04-09-2000

Patent document cited in search report		Publication date		Patent family member(s)		Publication date
EP 0712220	A	15-05-1996		JP 2986085 B		06-12-1999
				JP 8214004 A		20-08-1996
				US 5787071 A		28-07-1998
EP 0596159	A	11-05-1994		AU 4906393 A		19-05-1994
				CA 2102455 A		06-05-1994
				JP 6268664 A		22-09-1994